Gridline Graphs: A Review in Two Dimensions and an Extension to Higher Dimensions

Dale Peterson
Department of Mathematical Sciences
United States Air Force Academy
HQ USAFA/DFMS
2354 Fairchild Drive, Suite 6D2A
USAF Academy, CO 80840-6252
Dale.Peterson@usafa.af.mil

Abstract: *Gridline graphs* can be realized in the plane with vertices adjacent whenever they are on a common vertical or horizontal line. We review some applications and characterizations, e.g. they are line graphs of bipartite graphs, and provide practical $O(|V|^3)$ algorithms for some classical problems. We extend gridline graphs from the plane to higher dimensions. We characterize these graphs using a vertex coloring of the clique graph that corresponds to the conditions that, in the intersection graph of lines parallel to axes in $\mathbb{R}^n$, any cycle of four lines must remain in two dimensions and any path between two given lines must traverse the same two dimensions.

Keywords: Graph of a (0-1)-matrix; Adjacency graph; Checked graph; Gridline graph; Line graph of a bipartite graph; Clique graph; Perfect graph

1. Introduction and Background

A *gridline graph* is a graph that can be realized in the plane with vertices adjacent whenever they are on a common vertical or horizontal line. These graphs were introduced independently by Hedetniemi [6] who called them *graphs of (0,1)-matrices*; by Cook, Acharya, Devadas, and Mishra [2] who called them *adjacency graphs*; by Gurvich and Temkin [5] where the English translation called them *checked graphs*; by Beineke and Broere [1] who called them *rooks' graphs*; and by the author in [11].

Gridline graphs are used by Galvin [3] in his proof of the *Dinitz Conjecture*, which was that if each edge of the complete bipartite graph $K_{n,n}$ is assigned a list of n colors, then there is a proper edge coloring in which each edge is assigned a color from its list. Maffray [9] showed that a kernel in a certain gridline graph corresponds to a stable marriage system.

Edge and vertex colorings in general have application to many problems such as storage, scheduling, and fleet maintenance (see e.g. Roberts [14]). One such application is the *timetabling problem*, which is to assign course sections to periods so that no course has concurrent sections: Form a gridline graph with a vertex at (i, j) whenever teacher i is assigned to teach a section of class j, and then a proper vertex coloring corresponds to an assignment of periods.

Another application of gridline graphs is to *Rook polynomials* (see e.g. Roberts [14]), a combinatorial device for counting the number of ways objects can be placed on a board so that each object is in one of a set of allowable positions and no two objects are on a common line. This can be rephrased as counting the number of independent sets in a gridline graph.

One final application we mention is in *robotics*. Linear movement is easy for robots but turns are difficult. If movement is restricted to the vertical and horizontal and turns are allowable only at certain points, then a shortest path in a gridline graph gives the number of turns required. Allowing motion in a third dimension can be modeled by a 3-dimensional gridline graph.

In addition to their applications, gridline graphs are of interest because they are line graphs of bipartite graphs, which constitute an important class of perfect graphs. As such they are considered, for example, by Maffray and Preissmann [10] who studied induced subgraph vertex orderings for a type of sequential coloring, and by Hoffman [7,8] who considered spectra and other properties. For other references to these graphs, see the introductory paragraph in Hedetniemi [6].

We review some characterizations of gridline graphs, observe that they are perfect, and give practical polynomial-time algorithms for five classical problems. We extend gridline graphs from the plane to higher dimensions and characterize these graphs; these characterizations involve a vertex coloring that has a natural geometrical interpretation.

In this paper, a graph $G = (V, E)$ is undirected and has no multiple edges or loops. With a common abuse of language, we often refer to a vertex or edge as being in a graph G, and write, for example, $v \in G$ or $uw \in G$ instead of $v \in V(G)$ or $uw \in E(G)$. The cardinality of V (and, consequently, of E) is finite or denumerable. We take a *clique* to be a *maximal* complete subgraph. The *clique graph* $K(G)$ of a graph G has as its vertex set the cliques of G, with two vertices adjacent whenever they have some vertex of G in common. Cliques are typically denoted using capital letters. Hence, vertices of clique graphs are denoted by capital rather than small letters. The *line graph* $L(G)$ of a graph G has as its vertex set the edge-set of G, and vertices of $L(G)$ are adjacent whenever, as edges of G, they are adjacent. If G and H are graphs then G is *H-free* means that no induced subgraph of G is isomorphic to H. Also, a coloring is always proper.

2. Two-Dimensional Gridline Graphs

A *gridline graph* is a graph G that is isomorphic to some graph $\overline{G}$ – called a *realization* of G – whose vertices are located in $\mathbb{R}^2$, no two vertices have the same coordinates, and (distinct) vertices at (x, y) and (x', y') are adjacent whenever $x = x'$ or $y = y'$. That is, G can be realized in the plane such that no two vertices are colocated and two vertices are adjacent whenever they are on a common vertical or horizontal line. Any gridline graph can be realized in $\mathbf{N}^2$ since $V(G)$ is countable – hence the name. The term *line*, in the context of a gridline graph realization, always refers to a vertical or horizontal line.

Observe that, in a gridline graph, any two cliques contain at most one common vertex. The following lemma characterizes this property. It has been used often before (e.g. Tucker [15]), and its simple proof is omitted here (see Peterson [11,12]). A *diamond* is a K_4 minus one edge.

Lemma 2.1: *A graph is diamond-free if and only if any two cliques intersect in at most one vertex, which holds if and only if each edge lies in exactly one clique.* ∎

A *claw* is a $K_{1,3}$. A *hole* (or *n-hole*) of a graph G is an induced n-cycle where $n \geq 4$. An *odd hole* is a hole of odd length. Theorem 2.2 below characterizes gridline graphs. These characterizations were first due to Hedetniemi [6] and were independently noted by Peterson [12] and Beineke and Broere [1]. We briefly sketch a proof since some of its ideas prelude techniques used later in the paper.

Theorem 2.2: *For any graph G, the following statements are equivalent:*

(a) G is a gridline graph.

(b) G is the line graph of a bipartite graph.

(c) G is diamond-free and $K(G)$ is bipartite.

(d) G is diamond- and claw-free and has no odd hole.

Sketch of Proof: (a)⇒(d): None of the graphs in (d) can be (an induced subgraph of) a gridline graph. (d)⇒(c): If G has no odd hole and $K(G)$ is not bipartite, then it is straightforward to show, by contradiction, that a smallest odd cycle of $K(G)$ is a triangle. It is then easy to show (e.g. using Lemma 4.3 in Section 4) that G contains an induced claw or diamond. (c)⇒(a): Construct a gridline graph by placing cliques of G having one color on vertical lines and cliques having the other color on horizontal lines. By Lemma 2.1, no two vertices are colocated. (b)⇔(a): Given a bipartite graph H, construct the 0-1 matrix $\mathbf{A} = (a_{ij})$ where the rows and columns correspond to the two partite sets with $a_{ij} = 1$ iff the vertices corresponding to row i and column j are adjacent. Construct the gridline graph with a vertex at (i, j) iff $a_{ij} = 1$. It is immediate that $L(H)$ and G are isomorphic. Since a gridline graph can be realized with vertices only at positive integral points, the argument works in reverse. (The most concise full proof available is in Peterson [12].) ∎

Several facts about gridline graphs are immediate.

(1) In a realization, we can identify each clique with a unique line. (If a clique is an isolated vertex, then either line containing the vertex can be chosen to represent the clique.)

(2) A gridline graph does not have a unique realization.

(3) Any induced subgraph of a gridline graph is also a gridline graph, but this need not be true for a *partial* (not necessarily induced) subgraph – take a 5-cycle with and then without a chord.

(4) The set of *minimal forbidden induced subgraphs* – that is, the set $\{G : G$ is not a gridline graph but any induced subgraph is$\}$ – is not finite. For example, any odd hole is a minimal forbidden induced subgraph.

(5) If G is a gridline graph then we can construct a bipartite graph H such that $L(H) = G$ by first constructing $K(G)$, then (a) converting each isolated vertex into a component of two vertices and (b) adding a vertex of degree one for each vertex in G that is in only one clique.

(6) Gridline graphs are *perfect.* A graph is perfect whenever $\omega(G') = \chi(G')$ for every induced subgraph G' of G, where ω is the *clique number* (size of a maximum clique) and χ is the chromatic number. (In this definition and the remainder of the section, we assume that G is finite.) Perfectness of gridline graphs was proved by Hedetniemi [6] and independently by Gurvich and Temkin [5] and Peterson [11,12]. It also follows from the result by Tucker [15] that a diamond-free graph with no odd hole is perfect.

We conclude this section by noting that, for gridline graphs, five classical graph theory problems take time polynomial in the number of vertices (and edges). The emphasis below is to demonstrate polynomiality and simplicity of the algorithms. (Problems (1), (4), and (5) were shown to be polynomial for all perfect graphs by Grötschel, Lovász, and Schrijver [4]. Their approach is based on the ellipsoid method and uses a polynomial-time separation algorithm, and they do not recommend it for practical use.)

A simple algorithm for finding a nontrivial clique is to permute vertices of the adjacency matrix (that is, permute lines of the matrix) to obtain a maximal square of 1's in the upper-left corner, excluding the diagonal entries. If G is a gridline graph then, to find other cliques, remove these 1's, since each edge is in no other clique. If some edge of G is in two cliques – that is, G is not diamond-free – then, after finding some clique, say with r vertices, there will be another line in the adjacency matrix with at least two 1's among the first r positions.

To see the complexity of the above algorithm, first observe that if each vertex is in at most two cliques then the sum of the clique sizes is at most $2n$, where $|V| = n$; it then follows that there are at most n cliques. In finding another vertex in a clique, where k vertices have already been identified as being in the clique, the first k entries in at most $n - k$ lines are checked. Since $k(n-k) \leq n^2/4$, the total time to find a clique with size r is at most $n^2r/4$. Summing for each clique, the time is $O(n^3)$. The complexity of testing whether G is diamond-free is also $O(n^3)$, since after each clique there are at most $r(n-r) \le n^2/4$ checks.

(1) Finding the clique number and a maximum clique (in fact, all cliques). Find all cliques of G (using, say, the algorithm given above).

(2) Recognition. Find all cliques of G and test whether G is diamond-free. Then, construct $K(G)$ and determine whether it is bipartite (using, say, a breadth-first search, which takes time $O(m')$ where $m' = |E(K(G))| \le n(n-1)/2$).

(3) Realization. Find all cliques and a bipartition of $K(G)$. Construct the bipartite graph H as described in the fifth fact following the proof of Theorem 2.2. Then construct the matrix $\mathbf{A}$ and its associated gridline graph as in the proof of Theorem 2.2.

(4) Finding the chromatic number and a minimum coloring (more generally, finding a k-coloring where $k \ge \chi$). The chromatic number χ is equal to ω, found in problem (1). Tucker [15] gives an $O(kn^2)$ algorithm for perfect diamond-free graphs, and Maffray and Preissmann [10] give an $O(mn)$ algorithm (where $|E| = m$) for a class of perfect graphs containing gridline graphs.

(5) Finding the independence number and a maximum independent set. Find a maximum set of independent 1's in the matrix $\mathbf{A}$ mentioned in problem (3) (using, say, a maximum flow algorithm).

Some of these problems can be rephrased in terms of a bipartite graph H of which G is the line graph – for example, a maximum independent set of G corresponds to a maximum matching of H. From problem (3), H can be constructed in time $O(n^3)$. There are good algorithms for bipartite graphs, e.g. Schrijver [15] gives an $O(\Delta m)$ algorithm, where Δ is the maximum degree, for edge-coloring a bipartite graph, which corresponds to $O(\omega n)$ for vertex-coloring a gridline graph.

3. Higher Dimensional Gridline Graphs

In this section we extend gridline graphs from the plane to higher dimensions and characterize these graphs. A *p-dimensional* (or, for brevity, p-d) *gridline graph*, where $p \in \mathbf{N}$, is a graph G that is isomorphic to some graph $\overline{G}$ whose vertices are located in $\mathbb{R}^p$, no two vertices have the same coordinates, and vertices at $\mathbf{x} = (x_1, \ldots, x_p)$ and $\mathbf{x}' = (x'_1, \ldots, x'_p)$

are adjacent whenever they differ in exactly one entry. That is, G can be realized in $\mathbb{R}^p$ such that no two vertices are colocated and two vertices are adjacent whenever they are on a common line that is parallel to some axis. Any p-d gridline graph can be realized in $\mathbf{N}^p$. Analogous to a 2-dimensional gridline graph, a *realization* is a graph $\overline{G}$ as given in the definition, and the term *line*, in the context of a realization, always refers to a line parallel to some axis in $\mathbb{R}^p$.

The p-d gridline graphs were independently considered by Beineke and Broere [1], who called them *p-dimensional rooks' graphs.*

Like (2-d) gridline graphs, any induced subgraph of a p-d gridline graph is also a p-d gridline graph. But, by the same example used for gridline graphs (a 5-cycle with and then without a chord), a partial subgraph need not be a q-d gridline graph for any $q \in \mathbf{N}$. Also, a p-d gridline graph does not have a unique realization – for example, hyperplanes perpendicular to an axis can be permuted. Unlike 2-d gridline graphs, however, a p-d gridline graph need not be a line graph for any graph: Line graphs are claw-free, but a claw is a 3-d gridline graph.

We begin with four lemmas. The proofs are straightforward (see Peterson [11,12]) and are omitted here. Recall that vertices of the clique graph $K(G)$ are denoted by capital rather than small letters.

Lemma 3.1: *Suppose G is a diamond-free graph and that $\{A_j : j \in J \subseteq \mathbf{N}\}$ is a set of cliques in G that induces a complete subgraph in $K(G)$. Then there exists a unique vertex $v \in G$ such that $A_i \cap A_j = \{v\}$ whenever $i \neq j$.* ∎

By Lemma 3.1, diamond-free graphs satisfy the *clique-Helly* property (Prisner [13]), which is that if a set of cliques pairwise intersect then their intersection is nonempty.

Lemma 3.2: *If G is a diamond-free graph then so is $K(G)$.* ∎

Lemma 3.3: *Suppose G is a diamond-free graph. Then G has a 5-hole if (and, in fact, only if) $K(G)$ has a 5-hole.* ∎

Lemma 3.4: *Suppose H is a diamond-free graph and Z is a cycle in H containing the vertex a. If the two vertices consecutive to a in Z are not adjacent, then a is in a hole whose vertices are in Z.* ∎

Lemma 3.4 will be used in the following way: Suppose that G is a diamond-free graph, in which case, by Lemma 3.2, $K(G)$ is also diamond-free. Suppose further that $K(G)$ contains a cycle Z with vertex A satisfying the lemma conditions for a, and that the vertices of $K(G)$ are colored so that all other vertices of Z have a color different from A. Then A is in a hole whose other vertices all have a color different from A.

We proceed to characterize p-d gridline graphs, following a definition. In a vertex-colored graph, the color c *separates* vertices u and v whenever c appears on the interior of every (u, v)-path. Observe that if u and v are in different components then they are separated by every color.

Proposition 3.5: *A graph G is a p-d gridline graph if and only if it is is diamond-free and $K(G)$ is p-colorable such that (a) no hole contains some color exactly once, and (b) every pair of vertices at distance greater than two is separated by at least two colors.*

Remark: The p-colorability in the proposition statement has a geometrical interpretation. Any (A, B)-path in $K(G)$ corresponds to a path of lines (cliques of G) between the two lines l_A and l_B corresponding to A and B, respectively. If we let colors correspond to dimensions (parallel lines), then (a) says that any cycle of lines from l_A to l_B and back to l_A – where the cycle is actually in the intersection graph of lines – must traverse each intermediate dimension at least twice. Condition (b) says that there are two dimensions such that any path of lines from non-intersecting lines l_A to l_B – except for a direct line (if one exists) between l_A and l_B – must traverse those two dimensions. Suppose, for example, that l_A and l_B intersect, but there is no point of G at the intersection. Then any path of lines between l_A and l_B must contain a line parallel to l_A and one parallel to l_B. Similar geometric arguments hold for the other three cases: Lines l_A and l_B are nonparallel and nonintersecting, are parallel and in a common plane that is perpendicular to some axis, or are parallel and in no common plane perpendicular to some axis.

Proposition 3.5 states that these necessary geometric conditions, together with being diamond-free, are also sufficient.

Proof: In the case $p = 1$, G is a p-d gridline graph iff G is complete, so $K(G)$ is a single vertex. It is immediate that this satisfies the 1-coloring condition. Conversely, the given 1-coloring implies that $K(G)$ is an independent set; condition (b) implies that $K(G)$ is a single vertex, since otherwise there is only one color to separate the vertices having infinite distance.

Suppose $p \geq 2$. To avoid confusion between vertices in G and vertices in $K(G)$, we refer to vertices in G as *points*.

Before showing the equivalence of the statements, we make some important observations. The first is similar to the first fact following the proof of Theorem 2.2: In a realization of a p-d gridline graph, we can identify each clique with a unique line (parallel to one of the axes). Each such line – that is, each clique – can be represented by a p-vector with a $*$ in the entry corresponding to the parallel axis, and real numbers in the other $p - 1$ entries.

Parallel lines have a $*$ in the same entry and do not intersect.

Next, two cliques A and B of G are adjacent in $K(G)$ iff there is a point of G agreeing with both A and B on their respective $p-1$ non-$*$ entries. Thus A and B agree on their common $p-2$ non-$*$ entries. Two facts follow:

(i) If there is a path in $K(G)$ containing no vertex whose p-vector has a $*$ in entry r, then every vertex in the path agrees on entry r.

(ii) If there is an (A, B)-path in $K(G)$, the interior of which contains no vertex whose p-vector has a $*$ in entry r, then there are points of G in A and B that agree on entry r.

Finally, when using (b), we need not be concerned whether A and B are connected in $K(G)$; if not, then their distance $d(A, B)$ is infinite and (b) holds, since A and B are separated by p colors, $p \geq 2$.

("only if") Suppose G is a p-gridline graph. It is diamond-free since any induced subgraph must itself be a p-d gridline graph.

We may assume that G is a realization. For $r = 1, \ldots, p$, color each clique of G having a $*$ in entry r with the color r. We assume this coloring in the proofs of (a) and (b) below.

Proof of (a): Suppose $K(G)$ contains an n-cycle Z, $n \geq 4$, containing some color exactly once; w.l.o.g. it is the color 1 at vertex A, and A has p-vector $(*, a_2, \ldots, a_p)$. By (i), all vertices in Z other than A agree on the first entry, say at a_1. Thus a point at $(a_1, a_2, \ldots, a_p)$ is responsible for linking A to the two vertices consecutive to A in Z. But then these two vertices are also linked by this point, so Z has a chord.

Proof of (b): Suppose A and B are separated at most by one color, w.l.o.g. color 1. We show that there are points in A and B that differ on at most the first entry. This proves (b): If these points are the same then $d(A, B) \leq 1$, and otherwise they are in a common clique that intersects A and B, so $d(A, B) \leq 2$. Suppose A has color s and B has color t. By (ii), A and B differ at most on entries $1, s$, and t. We consider three cases.

Case 1: $s = t = 1$. Then $A = B$ since they each have the same p-vector, which has a $*$ in the first entry.

Case 2: $|\{1, s, t\}| = 2$. W.l.o.g. $s \neq 1$. Then A and B differ at most on entries one and s. By (ii), there are points in A and B that agree on entry s, so they differ at most on the first entry.

Case 3: $1, s, t$ are distinct. By (ii), A has a point that agrees with B on entry s, and B has a point that agrees with A on entry t. Then these two points agree on entries s and t, so they differ at most on the first entry.

("if") We will construct a p-d gridline graph in $\mathbf{N}^p$ that is isomorphic to G. Label the points of G as $v_1, v_3, \ldots$ (odd indices). Color $K(G)$ with the colors $1, \ldots, p$ according to (a) and (b). For each j, $j = 1, \ldots, p$, remove from $K(G)$ the vertices colored j and label the components of this graph as $K_{j,2}, K_{j,4}, \ldots$ (even indices).

Let $\mathbf{N}_*^p$ be the set of p-vectors with a $*$ in exactly one entry and positive integers in the other $p-1$ entries. Define the mapping $g : V(K(G)) \to \mathbf{N}_*^p$ by:

$$g_j(A) = \begin{cases} * & \text{if } A \text{ has color } j \\ s & \text{if } A \text{ is in } K_{j,s} \end{cases}$$

That is, the j^{th} entry of A is a $*$ iff A has color j, and two cliques not colored j agree on entry j iff, in $K(G)$, they are not separated by color j. We show that this is injective. Suppose in contradiction that two cliques A and B have the same p-vector. Then they have the same color, w.l.o.g. color 1. Now, for $j > 1$, $g_j(A) = g_j(B)$ means that at most color 1 separates A and B in $K(G)$. By (b), $d(A, B) \leq 2$. Since A and B have the same color but $A \neq B$, $d(A, B) = 2$. Then there is a path ACB where C has another color, w.l.o.g. color 2. Since A and B agree on the second entry, there is another (A, B)-path in which color 2 does not appear. These two paths form a cycle in $K(G)$ in which the vertices A and B that are consecutive to C are not adjacent, and C is the only vertex in the cycle with the color 2. By (the statement following) Lemma 3.4, this violates (a).

Now define the mapping $f : V(G) \to \mathbf{N}^p$ by:

$$f_j(v_\alpha) = \begin{cases} g_j(A) & \text{if } v_\alpha \in A \text{ for some clique } A \text{ where } g_j(A) \neq *, \\ & \text{that is, where } A \text{ is not colored } j \\ \alpha & \text{otherwise} \end{cases}$$

That is, v_α inherits the entries from the cliques in which it is contained; if it is in only one clique then the unassigned entry obtains the value α, which is unique to v_α. The mapping f is well defined since, if v_α is in two cliques A and B then they are adjacent, and by the mapping g they agree on their common $p-2$ non-$*$ entries. Note that, by this argument, facts (i) and (ii) hold for G with respect to the mappings f and g. We now proceed to show that G and the realization defined by f are isomorphic.

Vertex correspondence: We must show that f is injective, that is, that no points in $f(V)$ are colocated in the realization. Suppose in contradiction that $v_\theta \neq v_\phi$ but $f(v_\theta) = f(v_\phi)$. Observe that v_θ is in at least two cliques (necessarily having distinct colors), since otherwise one of its entries is the unique value θ. Similarly, v_ϕ is in at least two cliques. We consider two cases.

Case 1: v_θ and v_ϕ are in some common clique A. W.l.o.g. A has color 1. Both v_θ and v_ϕ inherited the same first entry from cliques B and C, respectively. Observe that $B \neq C$, since otherwise A and B both contain v_θ and v_ϕ, which by Lemma 2.1 violates that G is diamond-free. By the mapping g, B and C obtained the same first entry because there is a (B, C)-path in $K(G)$ not containing color 1. This path together with A form a cycle Z in $K(G)$. Now B and C are not adjacent, since otherwise ABC is a triangle in $K(G)$, and by Lemma 3.1 $v_\theta = v_\phi$, violating our hypothesis. But A is the only vertex colored 1 in Z, and its consecutive vertices B and C in Z are not adjacent; by Lemma 3.4 this violates (a).

Case 2: v_θ and v_ϕ are in no common clique. Suppose $v_\theta \in A_1$ and $v_\phi \in A_2$. Since v_θ and v_ϕ inherited their entries from the cliques containing them, A_1 and A_2 agree on their non-$*$ entries. It follows that they must have different colors, since otherwise $A_1 = A_2$ by the injection g. W.l.o.g. $f(v_\theta) = f(v_\phi) = (a_1, a_2, \ldots, a_p)$, A_1 has color 1 and so $g(A_1) = (*, a_2, \ldots, a_p)$, and A_2 has color 2 and so $g(A_2) = (a_1, *, \ldots, a_p)$. We show that A_1 and A_2 violate (b).

We first show that $d(A_1, A_2) > 2$. Suppose that A_1 and A_2 intersect in G. Then there is a point w at their intersection; it must also have p-vector $(a_1, a_2, \ldots, a_p)$. Observe that w cannot be v_θ or v_ϕ, since otherwise v_θ and v_ϕ are adjacent, violating case 2 hypothesis. Now we can apply case 1 to w and either v_θ or v_ϕ, reaching a contradiction. Now suppose that A_1 and A_2 both intersect some other clique A_3. Then A_3 has a third color, w.l.o.g. color 3, so A_3 has p-vector $(a_1, a_2, *, \ldots, a_p)$. Hence A_3 intersects with both A_1 and A_2 at points each with p-vector $(a_1, a_2, a_3, \ldots, a_p)$. These two points are distinct, since otherwise A_1 and A_2 intersect and we can apply case 1 to these two points to obtain a contradiction. Thus $d(A_1, A_2) > 2$.

Now we show A_1 and A_2 are not separated by at least two colors (in fact, by any color). Recall that v_θ is in some clique B other than A_1, and B cannot have the color 1. See Figure 3.1.

Figure 3.1: Vertex correspondence, case 2

Then, by the mapping g, there is a (B, A_2)-path not containing color 1, and concatenating A_1 at the front of this path yields an (A_1, A_2)-path not containing color 1 in its interior. (Similarly, there is an (A_1, A_2)-path not containing color 2 in its interior.) For $j > 2$, by the mapping g, there is an (A_1, A_2)-path not containing color j. Thus A_1 and A_2 violate (b), concluding case 2 and the vertex correspondence.

Edge correspondence: We must show:

$$v_\theta \text{ and } v_\phi \text{ are adjacent in } G \iff f(v_\theta) \text{ and } f(v_\phi) \text{ differ in exactly one entry}$$

($\Rightarrow$) Since v_θ and v_ϕ are adjacent they are in some common clique, so $f(v_\theta)$ and $f(v_\phi)$ agree in at least $p-1$ of their entries. By the vertex correspondence, $f(v_\theta) \neq f(v_\phi)$.

($\Leftarrow$) Suppose that $f(v_\theta)$ and $f(v_\phi)$ differ in exactly one entry, w.l.o.g. the first entry. We can write $f(v_\theta) = (a_1, a_2, \ldots, a_p)$ and $f(v_\phi) = (b_1, a_2, \ldots, a_p)$, where $a_1 \neq b_1$. Now suppose in contradiction that v_θ and v_ϕ are not adjacent, that is, are in no common clique. Let

$$\begin{aligned} \mathcal{V}_\theta &= \{\text{cliques (vertices of } K(G)) \text{ containing } v_\theta\} \\ \mathcal{V}_\phi &= \{\text{cliques (vertices of } K(G)) \text{ containing } v_\phi\} \end{aligned}$$

By supposition, $\mathcal{V}_\theta$ and $\mathcal{V}_\phi$ are disjoint. Any two cliques from $\mathcal{V}_\theta \cup \mathcal{V}_\phi$ agree on the non-$*$ entries, except possibly on the first entry. We consider four cases.

Case 1: There is a color 1 vertex in both $\mathcal{V}_\theta$ and $\mathcal{V}_\phi$, say $A_1 \in \mathcal{V}_\theta$ and $A_1' \in \mathcal{V}_\phi$. But then A_1 and A_1' agree on all entries, so by the injection g, $A_1 = A_1'$.

Case 2: One of $\mathcal{V}_\theta$ and $\mathcal{V}_\phi$ contains only one vertex, and it does not have color 1. W.l.o.g. $\mathcal{V}_\phi$ contains only A_2, with color 2. Then $f_2(v_\phi) = \phi$. But $f_2(v_\theta) \neq \phi$ since the value ϕ is unique to v_ϕ; this violates our assumption.

Case 3: Exactly one of $\mathcal{V}_\theta$ and $\mathcal{V}_\phi$ contains a color 1 vertex, and the other contains more than one vertex. W.l.o.g. $\mathcal{V}_\theta$ contains A_1, and $\mathcal{V}_\phi$ contains A_2 and A_3, where A_j has color j. A_1 has the p-vector $(*, a_2, a_3, \ldots, a_p)$, A_2 has $(b_1, *, a_3, \ldots, a_p)$, and A_3 has $(b_1, a_2, *, \ldots, a_p)$. See Figure 3.2. We show that $d(A_1, A_2) > 2$ and that A_1 and A_2 are separated only by color 1, violating (b).

Figure 3.2: Edge correspondence, case 3

Suppose there is an (A_1, A_2)-path not containing color 1 in its interior. Then, by (ii), A_1 contains a point with p-vector $(b_1, a_2, a_3, \ldots, a_p)$. But, by the vertex correspondence, this is v_ϕ, so v_θ and v_ϕ are in a common clique, violating our hypothesis. Thus any (A_1, A_2)-path contains a color 1 vertex in its interior. Since this vertex cannot be adjacent to A_1, $d(A_1, A_2) > 2$.

Next we show that A_1 and A_2 are not separated by color j, $j > 1$. By the mapping g there is an (A_1, A_3)-path not containing color 2; concatenating A_2 on the end of this path gives an (A_1, A_2)-path not containing color 2 in its interior. For $j > 2$, by the mapping g, there is an (A_1, A_2)-path not containing color j.

Case 4: $\mathcal{V}_\theta$ and $\mathcal{V}_\phi$ each contain at least two vertices, and neither contains a color 1 vertex. Then there are $A_2 \in \mathcal{V}_\theta$ and $A_3 \in \mathcal{V}_\phi$ where A_2 and A_3 have different colors. W.l.o.g. A_2 has color 2, A_3 has color 3, $A_r \in \mathcal{V}_\theta$ and has color $r \neq 2$, and $A_s \in \mathcal{V}_\phi$ and has color $s \neq 3$. See Figure 3.3. We show that $d(A_2, A_3) > 2$ and that A_2 and A_3 are separated only by color 1, violating (b).

Figure 3.3: Edge correspondence, case 4

Now A_2 and A_3 differ on the first entry and so are separated by color 1. Hence $d(A_2, A_3) \neq 1$. Suppose $d(A_2, A_3) = 2$. Then there is a path $A_2 B A_3$ where B has color 1. This implies that B has p-vector $(*, a_2, a_3, \ldots, a_p)$. Then B intersects A_2 and A_3 at points whose p-vectors are $(a_1, a_2, a_3, \ldots, a_p)$ and $(b_1, a_2, a_3, \ldots, a_p)$, respectively. But, by the vertex correspondence, these are v_θ and v_ϕ, so these points are in a common clique, violating our hypothesis. Thus $d(A_1, A_2) > 2$.

Next we show that A_1 and A_2 are not separated by color j, $j > 1$. By the mapping g there is an (A_r, A_3)-path not containing color 2; concatenating A_2 on the front of this path gives an (A_2, A_3)-path not containing color 2 in its interior. Similarly, using (A_2, A_s), there is an (A_2, A_3)-path not containing color 3 in its interior. For $j > 3$, by the mapping g, there is an (A_2, A_3)-path not containing color j.

This completes the edge correspondence and the proof. ∎

Before proceeding, we make some observations about Proposition 3.5.
(1) If $p = 2$ then Proposition 3.5 reduces to the equivalence of (a) and (c) in Theorem 2.2.
(2) A (finite) p-d gridline graph need not be perfect when $p > 2$. For example, a 7-cycle is a 3-d gridline graph since its clique graph, which is also a 7-cycle, can be colored consecutively as 1,2,3,1,2,3,2.
(3) Like (2-d) gridline graphs, p-d gridline graphs ($p > 2$) have no finite set of minimal forbidden induced subgraphs. Take a graph G as shown in Figure 3.4.

Figure 3.4: G satisfies (a) but not (b)

Then $K(G)$ is 3-colorable according to (a); the only two such colorings (within isomorphism of colors) are shown. But vertices x and y violate (b) in either coloring. If $k \geq 4$ then it is straightforward (though tedious) to check that removing any vertex from G yields either a 2-d or 3-d gridline graph (see Peterson [12] for details).
(4) Both (a) and (b) are needed in the proposition. From the previous fact, there are graphs satisfying (a) but not (b). To see that (a) is necessary, take G as a 5-cycle. Then $K(G)$ is also a 5-cycle, so (b) is satisfied since no pair of vertices have distance greater than two. But (a) is not satisfied for any p. It is true, however, that (a) can be relaxed to include only 4- and 5-holes. The following lemma will allow us to do this.

Lemma 3.6: *Suppose H is a diamond-free graph and γ is a p-coloring of H in which every pair of vertices at distance greater than two is separated by at least two colors. Then, using the coloring γ, no hole contains some color exactly once if and only if there is no 5-hole and every 4-hole is colored with exactly two colors.*

Proof: ("only if") This is immediate.

("if") Suppose H has no 5-hole and every 4-hole is colored with exactly two colors. Then no 4- or 5-hole contains some color exactly once. We must show the result for any n-hole where $n \geq 6$. Suppose that a, b, c, d, e, f are consecutive vertices in an n-hole Z where $n \geq 6$, and suppose in contradiction that some color appears exactly once on Z, w.l.o.g. color 1 appears only at c.

There must be a vertex g adjacent to both b and e, since otherwise these two vertices violate the color separation: Color 1 must be one of the separating colors, but color 1 does not appear in the (b, e)-path on the part of Z avoiding c. Then $bcdeg$ is a 5-cycle. It must have a chord; the possibilities are gc and gd. If both are present then $\{b, c, d, g\}$ induces a diamond. Edge gc alone implies that $cdeg$ is a 4-hole and would force e to be colored 1, violating our hypothesis. Thus $gd \in H$ and $gc \notin H$. Then $bcdg$ is a 4-hole, so b and d have the same color, w.l.o.g. color 2, and g is colored 1.

Now there must be a vertex h adjacent to both a and e, since otherwise color 1 does not separate them using the part of Z avoiding c. There are two cases.

Case 1: $g = h$. There must be a vertex i adjacent to b and f, since otherwise color 1 must be one of the separating colors by path $bgef$, but color 1 does not appear on the part of Z avoiding c. Now $i \neq g$, since otherwise $fg \in H$ and $\{d, e, f, g\}$ induces a diamond. Also $i \neq a$, since otherwise $af \in H$ and $agef$ is a 4-hole, which forces f to be colored 1 and violates our hypothesis. Then $bgefi$ is a 5-cycle. Possible chords are fg, ei, and gi. Edge fg is not in H as mentioned above. Edges ei and gi together imply that $\{b, g, e, i\}$ induces a diamond. Edge ei alone implies that $bgei$ is a 4-hole, which forces e to be colored 2 and violates a proper coloring. Edge gi alone implies that $efig$ is a 4-hole, which forces f to be colored 1 and violates our hypothesis. Thus there is no chord, so $bgefi$ is a 5-hole, a contradiction.

Case 2: $g \neq h$. Then $abgeh$ is a 5-cycle. Possible chords are ag, bh, and gh. We need not consider ag, since this is case 1. Edges bh and gh together imply that $\{b, g, e, h\}$ induces a diamond. Edge bh alone implies that $bgeh$ is a 4-hole, which forces e to be colored 2 and violates a proper coloring. Edge gh alone implies that $abgh$ is a 4-hole, which forces a to be colored 1 and violates our hypothesis. Thus there is no chord, so $abgeh$ is a 5-hole, a contradiction. ∎

A *p-gridline coloring* of a graph is a p-coloring in which every 4-hole is colored with exactly two colors and every pair of vertices at distance greater than two is separated by at least two colors. A graph that admits a p-gridline coloring is said to be *p-gridline colorable.* We now proceed to the main result of the section.

Theorem 3.7: *A graph G is a p-d gridline graph if and only if it is diamond-free, has no 5-hole, and $K(G)$ is p-gridline colorable.*

Proof: We observe that the conditions of the theorem are equivalent to those of Proposition 3.5. Condition (a) of the proposition for 5-holes implies that G has no 5-hole, and (a) and (b) imply p-gridline colorability. Conversely, since G is diamond-free and has no 5-hole, Lemmas 3.2 and 3.3 imply that the same holds for $K(G)$. Applying Lemma 3.6 with $H = K(G)$ implies (a) and (b). ∎

Recall the remark following the statement of Proposition 3.5, which gave a geometrical interpretation for the conditions of the proposition. Now, condition (a) of the proposition, that no hole contains some color exactly once, is essentially replaced with the (apparently) weaker condition that every 4-hole is colored with exactly two colors. This is the obviously necessary condition that any 4-cycle of lines in a realization must remain in two dimensions.

4. Blow-ups of Gridline Graphs

This section extends the results of the two previous sections to include graphs in which, in a realization, vertices may be colocated – in which case they are adjacent. We call these p-d *gridline blow-up graphs.*

A *blow-up* is a complete subgraph whose vertices all have the same closed neighborhood – so named because it is sometimes constructed by blowing up a single vertex. In a realization, vertices that are colocated constitute a blow-up. The inverse concept is the *reduced graph of* G, which is obtained from the graph G by reducing each maximal blow-up to a single vertex. (A graph containing no blow-ups is said to be *reduced*, or *canonical.*) It is immediate that a graph is a p-d gridline blow-up graph iff its reduced graph is a p-d gridline graph. Using the concepts of blow-up and reduced graph, we can obtain characterizations of p-d gridline blow-up graphs that are analogous to our earlier results.

A multigraph may have multiple edges and loops (though every multigraph in this section is bipartite and thus has no loop), and its vertex set and edge-multiset are each finite or denumerable. Repeated elements in the edge-multiset are distinguished as vertices in the line graph of a multigraph. If G is a graph and $\mathcal{G}$ is a set of graphs, then G is $\mathcal{G}$*-free* means no induced subgraph of G is isomorphic to any graph of $\mathcal{G}$.

The characterizations of p-d gridline blow-up graphs involve the graphs shown in Figure 4.1. They are a *4-fan* (also called a *gem*), a *4-wheel*, and a *stingray.*

Figure 4.1: 4-fan, 4-wheel, and stingray

Let $\mathcal{F}'$ consist of a claw, 4-fan, and 4-wheel, and let $\mathcal{F}$ be $\mathcal{F}'$ together with all cycles of odd length. Let $\mathcal{H}'$ consist of a stingray, 4-fan, and 4-wheel, and $\mathcal{H}$ consist of $\mathcal{H}'$ together with a 5-cycle. The two main results of this section are below.

Theorem 4.1: *For any graph G, the following statements are equivalent:*

(a) G is a (2-d) gridline blow-up graph.

(b) G is the line graph of a bipartite multigraph.

(c) $K(G)$ is bipartite.

(d) G is $\mathcal{F}$-free. $\blacksquare$

Theorem 4.2: *A graph G is a p-d gridline blow-up graph if and only if it is $\mathcal{H}$-free and*

$K(G)$ is p-gridline colorable. ■

Using the two lemmas below, these theorems are straightforward extensions of Theorem 2.2 and Proposition 3.5; the proofs are omitted. *Lemma 4.3* characterizes the property of a 2-d gridline blow-up graph that any vertex is in at most two cliques; the proof of Theorem 4.1 uses this lemma in place of Lemma 2.1. *Lemma 4.4* characterizes the property of a p-d gridline blow-up graph that vertices in two common cliques must be in a common blow-up; the proof of Theorem 4.2 relies on this lemma. The simple proof of Lemma 4.4 is omitted.

Lemma 4.3: *For any graph G, the following statements are equivalent:*

(a) G is $\mathcal{F}'$-free.

(b) No vertex of G is in more than two cliques.

(c) $K(G)$ is triangle-free.

Proof: (b)$\Rightarrow$(c): If $K(G)$ contains a triangle $A_1A_2A_3$ such that $A_1 \cap A_2 \cap A_3 = \emptyset$, then take three vertices from $A_1 \cap A_2$, $A_1 \cap A_3$, and $A_2 \cap A_3$, respectively. These three vertices induce a triangle and hence are in some clique different from A_1, A_2, or A_3. Thus each of these three vertices is in at least three cliques.

(c)$\Rightarrow$(a): Observe that, for each graph of $\mathcal{F}'$, the vertex of maximum degree is in at least three cliques and thus the clique graph contains a triangle. Since the clique graph of an induced subgraph of G is a (partial) subgraph of $K(G)$ (see e.g. [6] Lemma 7 or [12] Lemma 2.5), the result follows.

(a)$\Rightarrow$(b) Suppose vertex $v \in G$ is in three cliques A_1, A_2, and A_3; we show that G is not $\mathcal{F}'$-free. By maximality, there exists a vertex $v_1 \in A_1 \setminus A_2$, and there is a vertex $v_2 \in A_2 \setminus A_1$ such that $v_1v_2 \notin G$. Observe that $v \neq v_1, v_2$, and the pair $\{v_1, v_2\}$ is in no common clique. We consider two cases.

Case 1: Exactly one of $v_1, v_2 \in A_3$. W.l.o.g. $v_1 \notin A_3$ and $v_2 \in A_3$. Choose nonadjacent vertices $v_2' \in A_2 \setminus A_3$ and $v_3' \in A_3 \setminus A_2$. By our choices and case assumption, v, v_1, v_2, v_2', v_3' are distinct vertices since $v_2' \neq v\,(\in A_3), v_1\,(\notin A_2)$, $v_2\,(\in A_3)$ and $v_3' \neq v\,(\in A_2), v_1\,(\notin A_3)$, $v_2\,(\in A_2)$. Also, v is adjacent to the other four chosen vertices, and $v_2v_2', v_2v_3' \in G$. Then either $\{v, v_1, v_2', v_3'\}$ induces a claw or $\{v, v_1, v_2', v_2, v_3'\}$ induces a 4-fan or 4-wheel.

Case 2: Neither v_1 nor v_2 is in A_3. Choose $v_3' \in A_3 \setminus A_1$ not adjacent to v_1, and $v_3'' \in A_3 \setminus A_2$ not adjacent to v_2. If $v_3' = v_3''$ then $\{v, v_1, v_2, v_3'\}$ induces a claw. Otherwise, like case 1, we have that v, v_1, v_2, v_3', v_3'' are distinct vertices, v is adjacent to the other four chosen vertices, and $v_3'v_3'' \in G$. If $v_1v_3'' \notin G$ then $\{v, v_1, v_2, v_3''\}$ induces a claw; similarly if $v_2v_3' \notin G$. If $v_1v_3'', v_2v_3' \in G$ then $\{v, v_1, v_3', v_3'', v_2\}$ induces a 4-fan. ■

Lemma 4.4: *A graph G is $\mathcal{H}'$-free if and only if any vertices in more than one common clique are in a common blow-up.* ∎

Lemmas 3.1 through 3.4 can now be extended to account for blow-ups. The supposition of Lemmas 3.1 through 3.3 that G be diamond-free is replaced by G being $\mathcal{H}'$-free. This is because (by Lemmas 2.1 and 4.4) the reduced graph of a graph that is $\mathcal{H}'$-free is diamond-free, a graph and its reduced graph have (within isomorphism) the same clique graph, and (for Lemma 3.3) a graph has a 5-hole if and only if its reduced graph has a 5-hole. The single vertex v in the implication of Lemma 3.1 becomes a blow-up. The graph G in the statement following Lemma 3.4 can be changed from being diamond-free to being $\mathcal{H}'$-free. Now, the modified Lemmas 3.1 and 3.4 extend the proof of Proposition 3.5 (and Theorem 3.7) to Theorem 4.2.

The facts about gridline graphs noted after the proof of Theorem 2.2 also apply to 2-d gridline blow-up graphs, with a modification to the fifth fact: When obtaining H from $K(G)$, multiple edges can be recovered from vertices in G that are in a common blow-up. Perfectness of gridline blow-up graphs follows from a result by Maffray [9] that implies that line graphs of bipartite multigraphs are perfect. Indeed, it is equivalent to König's theorem that the edge chromatic number χ' is equal to the maximum degree Δ in a bipartite multigraph.

The algorithms of Section 2 can be applied to any 2-d gridline blow-up graph G by first constructing the reduced graph; this can be done in $O(n^3)$ time by comparing pairs of lines (rows or columns) in the adjacency matrix. After finding the cliques in the reduced graph, blow up each reduced vertex to its original vertices.

5. Conclusion

We have reviewed characterizations of gridline graphs in terms of line graphs, clique graphs, and forbidden subgraphs. Perfectness of these graphs follows from several previous results. Simple polynomial algorithms are sketched for the maximum clique, recognition, realization, vertex-coloring, and maximum independent set problems. In particular, the first three of these can be done in time $O(|V|^3)$ (as can the fourth and fifth, by applying referenced algorithms to a bipartite graph). Gridline graphs are extended to higher dimensions and characterized in terms of forbidden subgraphs together with a coloring of the clique graph. This coloring corresponds to the conditions in higher dimension Cartesian space that (1) any cycle of four lines (where lines must be parallel to some axis) must remain in two dimensions and (2) given any two lines, there are two dimensions such that any path of lines from one line to the other – except for a path of one line, if it exists – must traverse those two dimensions.

Among the questions that remain are, for higher dimensional gridline graphs, the complexity of various problems such as recognition and realization. Also, the algorithms for (2-dimensional) gridline graphs given in Section 2 almost certainly can be refined.

Acknowledgements

I gratefully acknowledge the support by the Office of Naval Research, grant N00450-93-1-0133, to Rutgers University. This grant supported my doctoral dissertation on which this paper is based.

I thank Lowell Beineke, Vladimir Gurvich, and Erich Prisner for providing references on related work. I thank Endre Boros, Peter Hammer, and Motakuri Ramana for their remarks to me regarding gridline graphs. I thank Therese Biedl and my advisor Fred Roberts for their careful review and comments on the paper. I thank the referee for pointing out the reference by Hedeniemi [6] and for making many other useful suggestions.

References

[1] L. Beineke and I. Broere, Rooks' graphs, unpublished (1995).

[2] C. R. Cook, B. Acharya, B. Devadas, and V. Mishra, Adjacency graphs, *Congr. Numer.* **X**, *Proc. Fifth Southeastern Conf. on Combinatorics, Graph Theory and Computing* (1974) 317–331.

[3] F. Galvin, The list chromatic index of a bipartite multigraph, *J. Combin. Theory Ser. B* **63** (1995) 153–158.

[4] M. Grötschel, L. Lovász, and A. Schrijver, Polynomial algorithms for perfect graphs, *Annals Discr. Math.* **21** (1984) 325–356.

[5] V. A. Gurvich and M. A. Temkin, Checked perfect graphs, *Soviet Math. Dokl.* **326** (1992) 227–232; translated into English, *Russian Acad. Sci. Dokl. Math.* **46** (1993) 248–253.

[6] S. T. Hedetniemi, Graphs of (0,1)-matrices, in: M. Capobianco, J. B. Frechen and M. Krolik, eds., **Recent Trends in Graph Theory (Proc. Conf., New York, 1970)**,

Lecture Notes in Mathematics **186** (Springer, Berlin, 1971) 157–171.

[7] A. J. Hoffman, On the line graph of the complete bipartite graph, *Ann. Math. Statist.* **34** (1964) 883–885.

[8] A. J. Hoffman, On eigenvalues and colorings of graphs, in: B. Harris, ed., **Graph Theory and its Applications; proceedings** (Academic Press, New York, 1969) 79–91.

[9] F. Maffray, Kernels in perfect line-graphs, *J. Combin. Theory Ser. B* **55** (1992) 1–8.

[10] F. Maffray and M. Preissmann, Sequential colorings and perfect graphs, *Discrete Appl. Math.* **94** (1999) 287 – 296.

[11] D. Peterson, Gridline graphs and higher dimensional extensions, *RUTCOR research report RRR 3-95*, Rutgers University, New Brunswick, NJ (1995).

[12] D. Peterson, Gridline graphs and the choice number of perfect line graphs, *RUTCOR dissertation RD #24*, Rutgers University, New Brunswick, NJ (1995).

[13] E. Prisner, Convergence of iterated clique graphs, *Discrete Math.* **103** (1992) 199–207.

[14] F. S. Roberts, **Applied Combinatorics** (Prentice-Hall, Englewood Cliffs, NJ, 1984).

[15] A. Schrijver, Bipartite edge coloring in $O(\Delta m)$ time, *SIAM J. Comput.* **28** (1998) 841–846.

[15] A. Tucker, Coloring perfect $(K_4 \setminus e)$-free graphs, *J. Combin. Theory Ser. B* **42** (1987) 313–318.

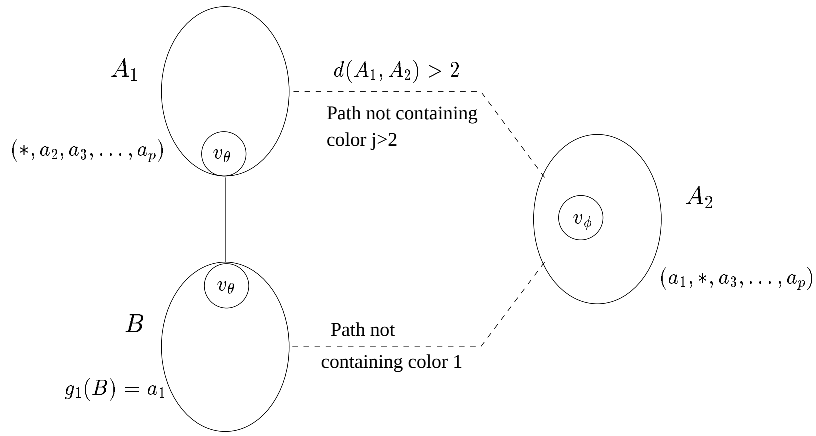

Figure 3.1: Vertex correspondence, case 2

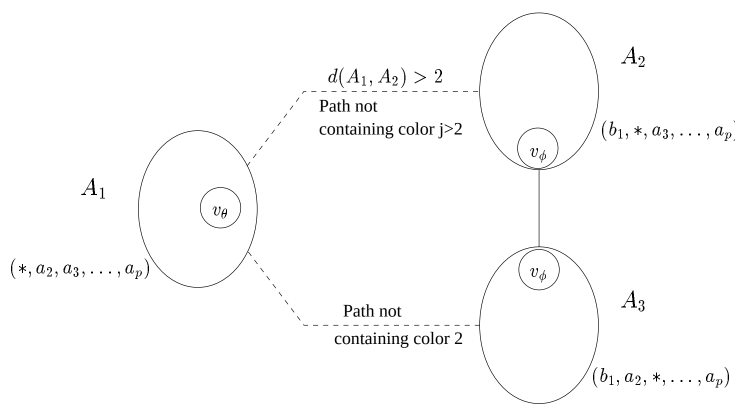

Figure 3.2: Edge correspondence, case 3

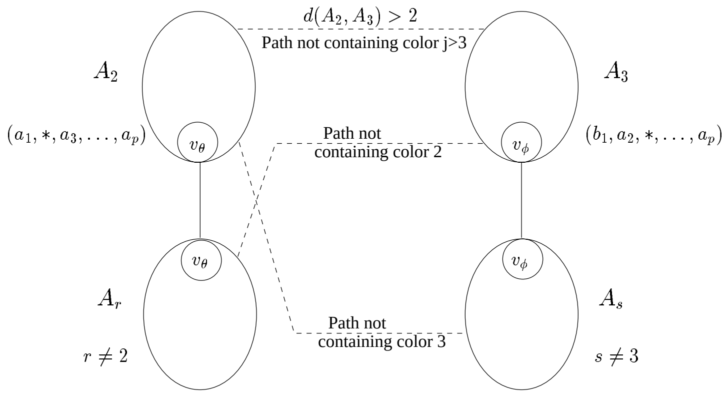

Figure 3.3: Edge correspondence, case 4

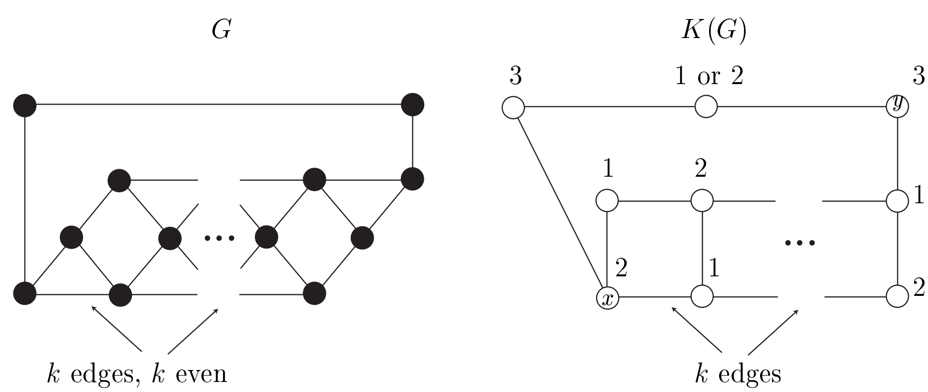

Figure 3.4: G satisfies (a) but not (b)

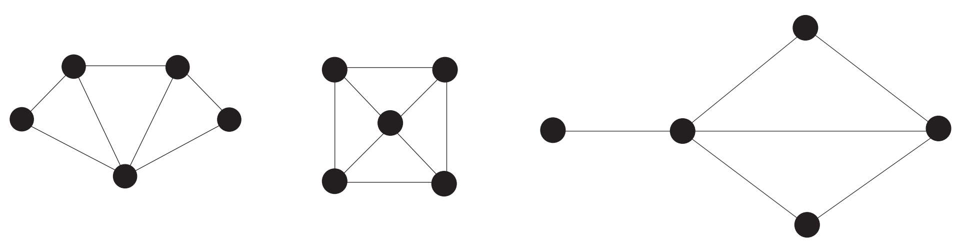

Figure 4.1: 4-fan, 4-wheel, and stingray